Optimal Event-Driven Multi-Agent Persistent Monitoring of a Finite Set of Targets

Nan Zhou[1], Xi Yu[2], Sean B. Andersson[1,2], and Christos G. Cassandras[1,3]
[1]Division of Systems Engineering, [2]Department of Mechanical Engineering , [3]Department of Electrical and Computer Engineering
Boston University, Boston, MA 02215, USA
E-mail:{nanzhou,xyu,sanderss,cgc}@bu.edu

arXiv:1603.04128v1 [math.OC] 14 Mar 2016

***Abstract*— We consider the problem of controlling the movement of multiple cooperating agents so as to minimize an uncertainty metric associated with a finite number of targets. In a one-dimensional mission space, we adopt an optimal control framework and show that the solution is reduced to a simpler parametric optimization problem: determining a sequence of locations where each agent may dwell for a finite amount of time and then switch direction. This amounts to a hybrid system which we analyze using Infinitesimal Perturbation Analysis (IPA) to obtain a complete on-line solution through an event-driven gradient-based algorithm which is also robust with respect to the uncertainty model used. The resulting controller depends on observing the events required to excite the gradient-based algorithm, which cannot be guaranteed. We solve this problem by proposing a new metric for the objective function which creates a potential field guaranteeing that gradient values are non-zero. This approach is compared to an alternative graph-based task scheduling algorithm for determining an optimal sequence of target visits. Simulation examples are included to demonstrate the proposed methods.**

I. Introduction

Systems consisting of cooperating mobile agents are often used to perform tasks such as coverage control [1], [2], surveillance, and environmental sampling. The *persistent monitoring* problem arises when agents must monitor a dynamically changing environment which cannot be fully covered by a stationary team of agents. Thus, persistent monitoring differs from traditional coverage tasks due to the perpetual need to cover a changing environment [3], [4]. A result of this exploration process is the eventual discovery of various "points of interest" which, once detected, become "targets" or "data sources" which need to be monitored. This setting arises in multiple application domains ranging from surveillance, environmental monitoring, and energy management [5], [6] down to nano-scale systems tasked to track fluorescent or magnetic particles for the study of dynamic processes in bio-molecular systems and in nano-medical research [7], [8]. In contrast to [3], [4] where *every* point in a mission space must be monitored, the problem we address here involves a *finite number* of targets (typically larger than the number of agents) which the agents must cooperatively monitor through periodic visits.

* The work of Cassandras and Zhou is supported in part by NSF under grants CNS-1239021, ECCS-1509084, and IIP-1430145, by AFOSR under grant FA9550-15-1-0471, and by ONR under grant N00014-09-1-1051. The work of Andersson and Yu is supported in part by the NSF through grant ECCS-1509084.

Each target may be viewed as a dynamic system in itself whose state is observed by agents equipped with sensing capabilities (e.g., cameras) and which are normally dependent upon their physical distance from the target. The objective of cooperative persistent monitoring in this case is to minimize an overall measure of uncertainty about the target states. This may be accomplished by assigning agents to specific targets or by designing motion trajectories through which agents reduce the uncertainty related to a target by periodically visiting it (and possibly remaining at the target for a finite amount of time). Viewed as an optimization problem, the goal is to jointly minimize some cost function that captures the desired features of the monitoring problem [9]. As long as the numbers of agents and targets is small, it is possible to identify sequences that yield a globally optimal solution; in general, however, this is a computationally complex procedure which does not scale well [10].

Rather than viewing this problem as a scheduling task which eventually falls within the class of traveling salesman or vehicle routing problems [11], in this paper we follow earlier work in [4] and introduce an optimal control framework whose objective is to control the movement of agents so as to collect information from targets (within agent sensing ranges) and ultimately minimize an average metric of uncertainty over all targets. An important difference between the persistent monitoring problem in previous work [3] and the current setting is that there is now a finite number of targets that agents need to monitor as opposed to every point in the mission space. In a one-dimensional mission space, we show that the optimal control problem can be reduced to a parametric optimization problem. In particular, every optimal agent trajectory is characterized by a finite number of points where the agent switches direction and by a dwelling time at each such point. As a result, the behavior of agents under optimal control is described by a hybrid system. This allows us to make use of Infinitesimal Perturbation Analysis (IPA) [12], [13] to determine on-line the gradient of the objective function with respect to these parameters and to obtain a (possibly local) optimal trajectory. Our approach exploits an inherent property of IPA which allows virtually arbitrary stochastic effects in modeling target uncertainty. Moreover, IPA's event-driven nature renders it *scalable* in the number of events in the system and not its state space.

A potential drawback of event-driven control methods is that they obviously depend on the events which "excite" the

controller being observable. However, this is not guaranteed under every feasible control: it is possible that no such events are excited, in which case the controller may be useless. The crucial events in our case are "target visits" and it is possible that such events may never occur for a large number of feasible agent trajectories which IPA uses to estimate a gradient on-line. At the heart of this problem is the fact that the objective function we define for a persistent monitoring problem has a non-zero cost metric associated with only a subset of the mission space centered around targets, while all other points have zero cost, since they are not "points of interest". This lack of event excitation is a serious problem in many trajectory planning and optimization tasks [14]–[16]. In this paper we solve this problem using a new cost metric introduced in [17] which creates a potential field guaranteeing that gradient values are generally non-zero throughout the mission space and ensures that all events are ultimately excited.

The rest of the paper is organized as follows. Section II formulates the optimal control problem and Section III presents a Hamiltonian analysis which characterizes the optimal solution in terms of two parameter vectors specifying switching points and associated dwelling times. In Section IV we provide a complete solution obtained through event-driven IPA gradient estimation, and solve the problem of potential lack of event excitation through a modified cost metric. Section V presents our graph-based scheduling approach and Section VI includes several simulation results.

II. Persistent Monitoring Problem Formulation

We consider N mobile agents moving in a one dimensional mission space $[0, L] \subset \mathbb{R}$. Let the position of the agents at time t be $s_j(t) \in [0, L]$, $j = 1, \ldots, N$, following the dynamics:

$$\dot{s}_j(t) = u_j(t) \tag{1}$$

i.e., we assume that the agent can control its direction and speed. Without loss of generality, after proper rescaling, we further assume that the speed is constrained by $|u_j(t)| \leq 1$, $j = 1, \ldots, N$. As will become clear, the agent dynamics in (1) can be replaced by a more general model of the form $\dot{s}_j(t) = g_j(s_n) + b_j u_j(t)$ without affecting the main results of our analysis. Finally, an additional constraint may be imposed if we assume that the agents are initially located so that $s_j(0) < s_{j+1}(0)$, $j = 1, \ldots, N-1$, and we wish to prevent them from subsequently crossing each other over all t:

$$s_j(t) - s_{j+1}(t) \leq 0 \tag{2}$$

The ability of an agent to sense its environment is modeled by a function $p_j(x, s_j)$ that measures the probability that an event at location $x \in [0, L]$ is detected by agent j. We assume that $p_j(x, s_j) = 1$ if $x = s_j$, and that $p_j(x, s_j)$ is monotonically nonincreasing in the distance $|x - s_j|$, thus capturing the reduced effectiveness of a sensor over its range which we consider to be finite and denoted by r_j. Therefore, we set $p_j(x, s_j) = 0$ when $|x - s_j| > r_j$. Although our analysis is not affected by the precise sensing model $p_j(x, s_j)$, we will limit ourselves to a linear decay model as follows:

$$p_j(x, s_j) = \max\{1 - \frac{|s_j - x|}{r_j}, 0\} \tag{3}$$

Unlike the persistent monitoring problem setting in [3], here we consider a known finite set of targets located at $x_i \in (0, L)$, $i = 1, \ldots, M$ (we assume $M > N$ to avoid uninteresting cases where there are at least as many agents as targets, in which case every target can be assigned to at least one agent). We can then set $p_j(x_i, s_j(t)) \equiv p_{ij}(s_j(t))$ to represent the effectiveness with which agent j can sense target i when located at $s_j(t)$. Accordingly, the joint probability that $x_i \in (0, L)$ is sensed by all N agents simultaneously (assuming detection independence) is

$$P_i(\mathbf{s}(t)) = 1 - \prod_{j=1}^{N} [1 - p_{ij}(s_j(t))] \tag{4}$$

where we set $\mathbf{s}(t) = [s_1(t), \ldots, s_N(t)]^{\mathrm{T}}$. Next, we define uncertainty functions $R_i(t)$ associated with targets $i = 1, \ldots, M$, so that they have the following properties: (i) $R_i(t)$ increases with a prespecified rate A_i if $P_i(\mathbf{s}(t)) = 0$ (we will later allow this to be a random process $\{A_i(t)\}$), (ii) $R_i(t)$ decreases with a fixed rate B_i if $P_i(\mathbf{s}(t)) = 1$ and (iii) $R_i(t) \geq 0$ for all t. It is then natural to model uncertainty dynamics associated with each target as follows:

$$\dot{R}_i(t) = \begin{cases} 0 & \text{if } R_i(t) = 0, A_i \leq B_i P_i(\mathbf{s}(t)) \\ A_i - B_i P_i(\mathbf{s}(t)) & \text{otherwise} \end{cases} \tag{5}$$

where we assume that initial conditions $R_i(0)$, $i = 1, \ldots, M$, are given and that $B_i > A_i > 0$ (thus, the uncertainty strictly decreases when there is perfect sensing $P_i(\mathbf{s}(t)) = 1$).

Our goal is to control the movement of the N agents through $u_j(t)$ in (1) so that the cumulative average uncertainty over all targets $i = 1, \ldots, M$ is minimized over a fixed time horizon T. Thus, setting $\mathbf{u}(t) = [u_1(t), \ldots, u_N(t)]$ we aim to solve the following optimal control problem **P1**:

$$\min_{\mathbf{u}(t)} \quad J = \frac{1}{T} \int_0^T \sum_{i=1}^{M} R_i(t) dt \tag{6}$$

subject to the agent dynamics (1), uncertainty dynamics (5), control constraint $|u_j(t)| \leq 1$, $t \in [0, T]$, and state constraints (2). Figure 1 is a polling model version for problem **P1** where each target is associated with a "virtual queue" where uncertainty accumulates with inflow rate A_i. The service rate of this queue is time-varying and given by $B_i P_i(\mathbf{s}(t))$, controllable through the agent position at time t. This interpretation is convenient for characterizing the *stability* of such a system over a mission time T: For each queue, we may require that $\int_0^T A_i < \int_0^T B_i P_i(\mathbf{s}(t)) dt$. Alternatively, we may require that each queue becomes empty at least once over $[0, T]$. Note that this analogy readily extends to two or three-dimensional settings.

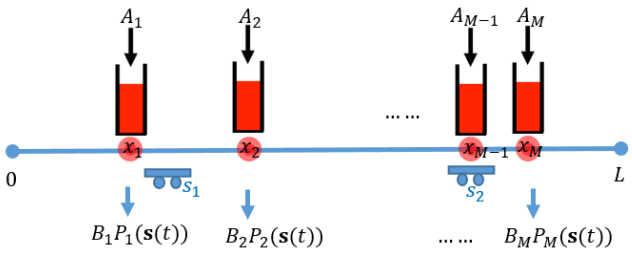

Fig. 1: A 1D polling model interpretation of problem **P1**

III. Optimal control solution

In this section, we derive properties of the optimal control solution of problem **P1** and show that it can be reduced to a parametric optimization problem. This will allow us to utilize an Infinitesimal Perturbation Analysis (IPA) gradient estimation approach [12] to find a complete optimal solution through a gradient-based algorithm. We begin by defining the state vector $\mathbf{x}(t) = [R_1(t), ...R_M(t), s_1(t)...s_N(t)]$ and associated costate vector $\lambda = [\lambda_1(t), ..., \lambda_M(t), \lambda_{s_1}(t), ..., \lambda_{s_N}(t)]$. As in [3], since the discontinuity in the dynamics of $R_i(t)$ in (5), the optimal state trajectory may contain a boundary arc when $R_i(t) = 0$ for some i; otherwise, the state evolves in an interior arc. Thus, we first analyze such an interior arc. Using (1) and (5), the Hamiltonian is

$$H(\mathbf{x}, \lambda, \mathbf{u}) = \sum_{i=1}^{M} R_i(t) + \sum_{i=1}^{M} \lambda_i(t)\dot{R}_i(t) + \sum_{j=1}^{N} \lambda_{s_j}(t)u_j(t) \tag{7}$$

The costate dynamics are

$$\dot{\lambda}_i(t) = -\frac{\partial H}{\partial R_i(t)} = -1, \quad \lambda_i(T) = 0 \tag{8}$$

$$\dot{\lambda}_{s_j}(t) = -\frac{\partial H}{\partial s_j(t)} = \sum_{i=1}^{M} \lambda_i(t) B_i \frac{\partial P_i(\mathbf{s}(t))}{\partial s_j(t)}, \quad \lambda_{s_j}(T) = 0 \tag{9}$$

Applying the Pontryagin Minimum Principle to (7) with $\mathbf{u}^\star(t)$, $t \in [0, T)$, denoting an optimal control, a necessary condition for optimality is

$$H(\mathbf{x}^\star, \lambda^\star, \mathbf{u}^\star) = \min_{u_j \in [-1,1],\ j=1,...,N} H(\mathbf{x}, \lambda, \mathbf{u}) \tag{10}$$

from which it immediately follows that

$$u_j^*(t) = \begin{cases} 1 & \text{if } \lambda_{s_j}(t) < 0 \\ -1 & \text{if } \lambda_{s_j}(t) > 0 \end{cases} \tag{11}$$

Note that there exists a possibility that $\lambda_{s_j}(t) = 0$ over some finite singular intervals [18], in which case $u_j^*(t)$ may take values in $\{-1, 0, 1\}$.

Similar to the case of the persistent monitoring problem studied in [3], the complete solution requires solving the costate equations (9), which in turn involves the determination of all points where $R_i(t) = 0$, $i = 1, \ldots, M$. This generally involves the solution of a two-point-boundary-value problem. However, we will next prove some structural properties of an optimal trajectory, based on which we show that it is fully characterized by a set of parameters, thus reducing the optimal control problem to a much simpler parametric optimization problem.

We begin by assuming that targets are ordered according to their location so that $x_1 < \cdots < x_M$. Let $r = \max_{j=1,...,N}\{r_j\}$ and $a = \max\{0, x_1 - r\}$, $b = \min\{L, x_M + r\}$. Thus, if $s_j(t) < x_1 - r$ or $s_j(t) > x_M + r$, then it follows from (3) that $p_{ij}(s_j(t)) = 0$ for all targets $i = 1, \ldots, M$. Clearly, this implies that the effective mission space is $[a, b]$, i.e.,

$$a \le s_j(t) \le b, \quad j = 1, \ldots, N \tag{12}$$

imposing an additional state constraint for **P1**. We will show next that on an optimal trajectory every agent is constrained to move within the interval $[x_1, x_M]$. This implies that every agent must switch its direction no later than reaching the first or last target (possibly after dwelling at the switching point for a finite time interval). To establish this and subsequent results, we will make a technical assumption that no two events altering the dynamics in this system can occur at the exact same time.

Assumption 1: Suppose that an agent switches direction at $\theta \in [a, b]$. For any $j = 1, \ldots, N$, $i = 1, \ldots, M$, $t \in (0, T)$, and any $\epsilon > 0$, if $s_j(t) = \theta$, $s_j(t - \epsilon) > \theta$ or if $s_j(t) = \theta$, $s_j(t - \epsilon) < \theta$, then either $R_i(\tau) > 0$ for all $\tau \in [t - \epsilon, t]$ or $R_i(\tau) = 0$ for all $\tau \in [t - \epsilon, t]$.

Proposition 1: In an optimal trajectory, $x_1 \le s_j^*(t) \le x_M$, $t \in [0, T]$, $j = 1, \ldots, N$.

Proof. We first prove that $s_j^*(t) \ge x_1$ for any agent j. Suppose that $s_j^*(t_0) = x_1$ and $u_j^*(t_0) = -1$. In view of (12), assume that agent j reaches a point $\theta \in [a, x_1)$ at time $t_1 > t_0$ where it switches direction; we will show that $\theta \notin [a, x_1)$ using a contradiction argument. There are two cases to consider:

Case 1: $\theta = a$. Assuming $s_j^*(t_1) = a$, we first show that $\lambda_{s_j}^*(t_1^-) = 0$ by a contradiction argument. If $\lambda_{s_j}^*(t_1^-) \neq 0$, recall that $u_j^*(t_1^-) = -1$, therefore $\lambda_{s_j}^*(t_1^-) > 0$ from (11). Since the constraint $a - s_j(t) \le 0$ is active, $\lambda_{s_j}^*(t)$ may experience a discontinuity so that

$$\lambda_{s_j}^*(t_1^-) = \lambda_{s_j}^*(t_1^+) - \pi_j \tag{13}$$

where $\pi_j \ge 0$ is a scalar multiplier associated with the constraint $a - s_j(t) \le 0$. It follows that $\lambda_{s_j}^*(t_1^+) = \lambda_{s_j}^*(t_1^-) + \pi_j > 0$. Since the Hamiltonian in (7) and the constraint $a - s_j(t) \le 0$ are not explicit functions of time, we have [18] $H^*(\mathbf{x}(t_1^-), \lambda(t_1^-), \mathbf{u}(t_1^-)) = H^*(\mathbf{x}(t_1^+), \lambda(t_1^+), \mathbf{u}(t_1^+))$ which, under Assumption 1, reduces to

$$\lambda_{s_j}^*(t_1^-)u_j^*(t_1^-) = \lambda_{s_j}^*(t_1^+)u_j^*(t_1^+) \tag{14}$$

Recall that $\lambda_{s_j}^*(t_1^-)u_j^*(t_1^-) < 0$. However, $u_j^*(t_1^+) \ge 0$ (since the agent switches control), therefore $\lambda_{s_j}^*(t_1^+)u_j^*(t_1^+) \ge 0$ which violates (14). This contradiction implies that $\lambda_{s_j}^*(t_1^-) = 0$. Recalling (4) and (9), we get $\dot{\lambda}_{s_j}^*(t_1^-) =$

$\sum_{i=1,R_i\neq 0}^{M}\lambda_i^*(t_1^-)\frac{B_i}{r_j}\prod_{d\neq j}[1-p_{id}(s_d^*(t_1^-))]$. Under Assumption 1, there exists $\delta>0$ such that during interval $(t_1-\delta,t_1)$, no $R_i(t)\geq 0$ becomes active, hence no $\lambda_i^*(t)$ encounters a jump for $i=1,\ldots,M$ and it follows from (8) that $\lambda_i^*(t)>0$. Moreover, $p_{id}(s_d^*(t))\neq 1$ for at least some $d\neq j$ since we have assumed that $M>N$. Thus, we have $\dot{\lambda}_{s_j}^*(t)>0$, for all $t\in(t_1-\delta,t_1)$. However, since agent j is approaching a, there exists some $\delta'<\delta$, such that $u_j^*(t)=-1$ for all $t\in(t_1-\delta',t_1)$, and $\lambda_{s_j}^*(t)\geq 0$. Thus for $t\in(t_1-\delta',t_1)$, we have $\lambda_{s_j}^*(t)\geq 0$ and $\dot{\lambda}_{s_j}^*(t)>0$. This contradicts the established fact that $\lambda_{s_j}^*(t_1^-)=0$. We conclude that $\theta\neq a$.

Case 2: $\theta\in(a,x_1)$. Assuming $s_j^*(t_1)=\theta$, we still have $u_j^*(t_1^-)=-1$, $u_j^*(t_1^+)\geq 0$. Since the Hamiltonian (7) is not an explicit function of time, we have $H^*(\mathbf{x}(t_1^-),\lambda(t_1^-),\mathbf{u}(t_1^-))=H^*(\mathbf{x}(t_1^+),\lambda(t_1^+),\mathbf{u}(t_1^+))$ which leads to (14) under Assumption 1. First, we assume $\lambda_{s_j}^*(t_1^-)\neq 0$. Since $u_j^*(t_1^-)<0$, we have $\lambda_{s_j}^*(t_1^-)>0$ and the left hand side of (14) gives $\lambda_{s_j}^*(t_1^-)u_j^*(t_1^-)<0$. On the other hand, in order to satisfy (14), we must have $u_j^*(t_1^+)>0$ and $\lambda_{s_j}^*(t_1^+)<0$. However, if $\lambda_{s_j}^*(t_1^-)>0$ and $\lambda_{s_j}^*(t_1^+)<0$, then either $\dot{\lambda}_{s_j}^*(t_1)<0$ and $\lambda_{s_j}^*(t_1)=0$, or $\lambda_{s_j}^*(t)$ experiences a discontinuity at t_1. We show that neither condition is feasible. The first one violates our assumption that $\lambda_{s_j}^*(t_1)\neq 0$, while the second one is not feasible since at $t=t_1$ the constraint $a-s_j(t)\leq 0$ is not active. This implies that $\lambda_{s_j}^*(t_1^-)=0$. Again, under Assumption 1, the same argument as in *Case 1* can be used to show that $\lambda_{s_j}^*(t)\geq 0$ and $\dot{\lambda}_{s_j}^*(t)>0$ for all $t\in(t_1-\delta',t_1)$. This contradicts the established fact that $\lambda_{s_j}^*(t_1^-)=0$ and we conclude that $\theta\notin(a,x_1)$.

Combining both cases, we conclude that $\theta\notin[0,x_1)$, which implies that $s_j^*(t)\geq x_1$. The same line of argument can be used to show that $s_j^*(t)\leq x_M$.■

Proposition 1, in conjunction with (11), leads to the conclusion that the optimal control consists of each agent moving with maximal speed in one direction until it reaches a point in the interval $[x_1,x_M]$ where it switches direction. However, the exclusion of the case $\lambda_{s_j}(t)=0$ allows the possibility of *singular arcs* along the optimal trajectory, defined as intervals $[t_1,t_2]$ such that $\lambda_{s_j}(t)=0$ for all $t\in[t_1,t_2]$ and $\lambda_{s_j}(t_1^-)\neq 0$, $\lambda_{s_j}(t_2^+)\neq 0$. The next result establishes the fact that we can exclude singular arcs from an agent's trajectory while this agent has no target in its sensing range.

Lemma 1: If $|s_j(t)-x_i|>r_j$ for any $i=1,\ldots,M$, then $u_j^*(t)\neq 0$.

Proof. We proceed with a contradiction argument. Suppose that $u_j^*(t)=0$ for $t\in[t_1,t_2]$ such that $|s_j^*(t_1)-x_i|>r_j$ for all $i=1,\ldots,M$ and that $u_j^*(t)\neq 0$ (without loss of generality, let $u_j^*(t)=1$) for $t>t_2$ so that $|s_j^*(t_3)-x_i|=r_j$ for some $i=1,\ldots,M$ and $|s_j^*(t_3+\Delta)-x_i|<r_j$ for $t_3+\Delta>t_3>t_2$. In other words, agent j eventually reaches a target i that it can sense at $t=t_3$. Assume that $u_j^*(t)$, $t\in[t_1,t_3+\Delta]$ is replaced by $u_j'(t)$ as follows: $u_j'(t)=1$ for $t\in[t_1,t_3+\Delta+t_1-t_2]$ and $u_j'(t)=0$ for $t\in(t_3+\Delta+t_1-t_2,t_3+\Delta]$. In other words, the agent moves to reach $s_j'(t_3+\Delta+t_1-t_2)=s_j^*(t_3+\Delta)$ and then stops. The two controls are thereafter identical. Then, referring to (6) we have $\int_{t_3+\Delta+t_1-t_2}^{t_3+\Delta}R_i'(t)dt\leq\int_{t_3+\Delta+t_1-t_2}^{t_3+\Delta}R_i^*(t)dt$ since under $u_j'(t)$ the agent may decrease $R_i(t)$ over $[t_3+\Delta+t_1-t_2,t_3]$ whereas under $u_j^*(t)$ this is impossible since $|s_j^*(t)-x_i|>r_j$ over this time interval. Since the cost in (6) is the same over $[0,t_3+\Delta+t_1-t_2)$ and $(t_3+\Delta,T]$, it follows that $u_j^*(t)=0$ when $|s_j(t)-x_i|>r_j$ cannot be optimal unless $u_j^*(t)=0$ for all $t\in[0,T]$, i.e., the agent never moves and never senses any target, in which case the cost under $u_j'(t)$ is still no higher than that under $u_j^*(t)$.■

Based on Lemma 1, we conclude that singular arcs in an agent's trajectory may occur only while it is sensing a target. Intuitively, this indicates that it may be optimal for an agent to stop moving and dwell in the vicinity of one or more targets that it can sense so as to decrease the associated uncertainty functions to an adequate level before it proceeds along the mission space. The next lemma establishes the fact that if the agent is visiting an *isolated* target and experiences a singular arc, then the corresponding optimal control is $u_j^*(t)=0$. An isolated target with position x_i is defined to be one that satisfies $|x_i-x_j|>2r$,for all $j\neq i$ where r was defined earlier as $r=\max_{j=1,\ldots,N}\{r_j\}$. Accordingly, the subset $I\sqsubseteq\{1,\ldots,M\}$ of isolated targets is defined as

$$I=\{i:|x_i-x_j|>2r, j\neq i\in\{1,...,M\}, r=\max_{j=1,\ldots,N}\{r_j\}\} \tag{15}$$

Lemma 2: Let $|s_j^*(t)-x_k|<r_j$ for some $j=1,\ldots,N$ and isolated target $k\in I$. If $\lambda_{s_j}^*(t)=0$, $t\in[t_1,t_2]$, then $u_j^*(t)=0$.

Proof. The proof is along the same line as Proposition III.3 in [3]. Assume that $\lambda_{s_j}^*(t)=0$ over a singular arc $[t_1,t_2]$. Let $H^*\equiv H(\mathbf{x}^*,\lambda^*,\mathbf{u}^*)$. Since the Hamiltonian along an optimal trajectory is a constant, we have $\frac{dH^*}{dt}=0$. Therefore, recalling (7),

$$\frac{dH^*}{dt}=\sum_{i=1}^{M}\left[\dot{R}_i^*(t)+\dot{\lambda}_i^*(t)\dot{R}_i^*(t)+\lambda_i^*(t)\ddot{R}_i^*(t)\right]+\sum_{j=1}^{N}\left[\dot{\lambda}_{s_j}^*(t)u_j^*(t)+\lambda_{s_j}^*(t)\dot{u}_j^*(t)\right]=0$$

and since $\dot{\lambda}_i^*(t)=-1$ from (8), this reduces to

$$\frac{dH^*}{dt}=\sum_{i=1}^{M}\lambda_i^*(t)\ddot{R}_i^*(t)+\sum_{j=1}^{N}\left[\dot{\lambda}_{s_j}^*(t)u_j^*(t)+\lambda_{s_j}^*(t)\dot{u}_j^*(t)\right]=0 \tag{16}$$

Define $S(t)=\{j|\lambda_{s_j}(t)=0,\dot{\lambda}_{s_j}(t)=0\}$ as the set of agents in singular arcs at t and $\bar{S}(t)$ as the set of all remaining agents. If $j\in S(t)$, then $\dot{\lambda}_{s_j}^*(t)u_j^*(t)+\lambda_{s_j}^*(t)\dot{u}_j^*(t)=0$. If $j\in\bar{S}(t)$, then $\lambda_{s_j}^*(t)\dot{u}_j^*(t)=0$ since $u_j^*(t)=\pm 1$ and $\dot{u}_j^*(t)=0$. Therefore, we rewrite (16) as

$$\frac{dH^*}{dt}=\sum_{i=1}^{M}\lambda_i^*(t)\ddot{R}_i^*(t)+\sum_{j\in\bar{S}(t)}\dot{\lambda}_{s_j}^*(t)u_j^*(t)=0 \tag{17}$$

Recalling (5), when $R_i(t) \neq 0$, we have $\dot{R}_i = A_i - B_i\Big(1 - \prod_{n=1}^{N}\Big[1 - p_{ij}\big(s_j(t)\big)\Big]\Big)$. Therefore,

$$\begin{aligned}\ddot{R}_i^*(t) =& \frac{d}{dt}\dot{R}_i^*(t) \\ =& -\sum_{j=1}^{N} u_j^*(t) B_i \frac{\partial p_{ij}(s_j^*(t))}{\partial s_j^*} \prod_{d\neq j}\Big[1 - p_{id}(s_d^*(t))\Big]\end{aligned} \tag{18}$$

Moreover, from (9), we have

$$\dot{\lambda}_{s_j}^*(t) = \sum_{i=1, R_i\neq 0}^{M} \lambda_i^*(t) B_i \frac{\partial p_{ij}(s_j^*(t))}{\partial s_j^*} \prod_{d\neq j}\Big[1 - p_{id}(s_d^*(t))\Big] \tag{19}$$

Combining (17)-(19), we get

$$\begin{aligned}&\frac{dH^*}{dt} = \\ &-\sum_{\substack{i=1\\R_i\neq 0}}^{M}\sum_{j=1}^{N} u_j^*(t)\lambda_i^*(t) B_i \frac{\partial p_{ij}(s_j^*(t))}{\partial s_j^*} \prod_{d\neq j}\Big[1 - p_{id}(s_d^*(t))\Big] \\ &+\sum_{j\in\bar{S}(t)}\sum_{\substack{i=1\\R_i\neq 0}}^{M} u_j^*(t)\lambda_i^*(t) B_i \frac{\partial p_{ij}(s_j^*(t))}{\partial s_j^*} \prod_{d\neq j}\Big[1 - p_{id}(s_d^*(t))\Big] \\ &=\sum_{j\in S(t)}\sum_{\substack{i=1\\R_i\neq 0}}^{M} u_j^*(t)\lambda_i^*(t) B_i \frac{\partial p_{ij}(s_j^*(t))}{\partial s_j^*} \prod_{d\neq j}\Big[1 - p_{id}(s_d^*(t))\Big] \\ &=0\end{aligned} \tag{20}$$

Since we have assumed that $|s_j^*(t) - x_k| < r_j$ and k is an isolated target, it follows that $p_{kj}(s_j^*(t)) \neq 0$ and $p_{ij}(s_j(t)) = 0$ if $i \neq k$. Therefore, $\frac{\partial p_{kj}(s_j^*(t))}{\partial s_j^*} \neq 0$ and $\frac{\partial p_{ij}(s_j^*(t))}{\partial s_j^*} = 0$ for all $i \neq k$ and (20) reduces to

$$\sum_{j\in S(t)} u_j^*(t)\lambda_k^*(t) B_i \frac{\partial p_{kj}(s_j^*(t))}{\partial s_j^*} \prod_{d\neq j}\Big[1 - p_{kd}(s_d^*(t))\Big] = 0 \tag{21}$$

Observe that, from (8), $\lambda_i(t) > 0$ when $R_i(t) \neq 0, t < T$. In addition $B_i > 0$ and $\prod_{d\neq j}[1 - p_{kd}(s_d^*(t))] \neq 0$. Therefore, to satisfy (21) for all $t \in [t_1, t_2]$, we must have $u_j^*(t) = 0$, for all $j \in S(t)$.■

We can further establish the fact that if an agent j experiences a singular arc while sensing an isolated target k, then the optimal point to stop is such that $s_j^*(t) = x_k$.

Proposition 2: Let $|s_j^*(t) - x_k| < r_j$ for some $j = 1, \ldots, N$ and isolated target $k \in I$. If $\lambda_{s_j}^*(t) = 0$, $t \in [t_1, t_2]$, and $u_j^*(t_1^-) = u_j^*(t_2^+)$, then $s_j^*(t) = x_k$, $t \in [t_1, t_2]$.

Proof. By Lemma 2, we know that $u_j^*(t) = 0$, $t \in [t_1, t_2]$. We use a contradiction argument similar to the one used in Lemma 1 to show that $s_j^*(t) = x_k$, $t \in [t_1, t_2]$. Suppose that $u_j^*(t_1^-) = 1$ (without loss of generality) and that $s_j^*(t) = x_k - \Delta < x_k$. Note that at the end of the singular arc $u_j^*(t_2^+) = 1$ since $u_j^*(t_1^-) = u_j^*(t_2^+)$. This implies that $s_j^*(t_2 + \Delta) = x_k$. Assume that $u_j^*(t)$, $t \in [t_1, t_2 + \Delta]$ is replaced by $u_j'(t)$ as follows: $u_j'(t) = 1$ for $t \in [t_1, t_1 + \Delta]$ and $u_j'(t) = 0$ for $t \in (t_1 + \Delta, t_2 + \Delta]$. In other words, the agent moves to reach $s_j'(t_1 + \Delta) = s_j^*(t_2 + \Delta) = x_k$ and then stops. The two controls are thereafter identical. Then, referring to (6) we have $\int_{t_1}^{t_2+\Delta} R_i'(t)dt < \int_{t_1}^{t_2+\Delta} R_i^*(t)dt$ since $\dot{R}_i^*(t) < \dot{R}_i'(t)$ due to (5) and the fact that $p_{kj}(s_j(t))$ is monotonically decreasing in $|s_j(t) - x_k|$. Since the cost in (6) is the same over $[0, t_1)$ and $(t_2 + \Delta, T]$, it follows that $s_j^*(t) = x_k - \Delta$ cannot be optimal. The same argument holds for any $\Delta > 0$, leading to the conclusion that $s_j^*(t) = x_k$, $t \in [t_1, t_2]$. A similar argument also applies to the case $s_j^*(t) = x_k + \Delta > x_k$.■

Finally, we consider the case with the state constraint (2). We can then prove that this constraint is never active on an optimal trajectory, i.e., agents reverse their directions before making contact with any other agent.

Proposition 3: Under the constraint $s_j(t) \leq s_{j+1}(t)$, on an optimal trajectory, $s_j(t) \neq s_{j+1}(t)$ for all $t \in (0, T)$, $j = 1...N$.

Proof. The proof is almost identical to that of Proposition III.4 in [3] and is, therefore, omitted.■

The above analysis, including Propositions 1-3, fully characterize the structure of the optimal control as consisting of intervals in $[0, T]$ where $u_j^*(t) \in \{-1, 0, 1\}$ depending entirely on the sign of $\lambda_{s_j}(t)$. Based on this analysis, we can parameterize **P1** so that the cost in (6) depends on a set of (i) *switching points* where an agent switches its control from $u_j(t) = \pm 1$ to ∓ 1 or possibly 0, and (ii) *dwelling times* if an agent switches from $u_j(t) = \pm 1$ to 0. In other words, the optimal trajectory of each agent j is totally characterized by two parameter vectors: switching points $\boldsymbol{\theta}_j = [\theta_{j1}, \theta_{j2}...\theta_{j\Gamma}]$ and dwelling times $\boldsymbol{\omega}_j = [\omega_{j1}, \omega_{j2}...\omega_{j\Gamma'}]$ where Γ and Γ' are prior parameters depending on the given time horizon. This defines a hybrid system with state dynamics (1), (5). The dynamics remain unchanged in between events that cause them to change, i.e., the points $\theta_{j1}, \ldots, \theta_{j\Gamma}$ above and instants when $R_i(t)$ switches from > 0 to 0 or vice versa. Therefore, the overall cost function (6) can be parametrically expressed as $J(\boldsymbol{\theta}, \boldsymbol{\omega})$ and rewritten as the sum of costs over corresponding interevent intervals over a given time horizon:

$$J(\boldsymbol{\theta}, \boldsymbol{\omega}) = \frac{1}{T}\sum_{k=0}^{K}\int_{\tau_k(\boldsymbol{\theta},\boldsymbol{\omega})}^{\tau_{k+1}(\boldsymbol{\theta},\boldsymbol{\omega})}\sum_{i=1}^{M} R_i(t)dt \tag{22}$$

where τ_k is the k-th event time. This will allow us to apply IPA to determine a gradient $\nabla J(\boldsymbol{\theta}, \boldsymbol{\omega})$ with respect to these parameters and apply any standard gradient-based optimization algorithm to obtain a (locally) optimal solution.

IV. Infinitesimal Perturbation Analysis

As concluded in the previous section, optimal agent trajectories may be selected from the family $\{\boldsymbol{s}(\boldsymbol{\theta}, \boldsymbol{\omega}, t, \boldsymbol{s}_0)\}$ with parameter vectors $\boldsymbol{\theta}$ and $\boldsymbol{\omega}$ and a given initial condition $\boldsymbol{s}_0$. Along these trajectories, agents are subject to dynamics (1) and targets are subject to (5). An *event* (e.g., an agent stopping at some target x_i) occurring at time $\tau_k(\boldsymbol{\theta}, \boldsymbol{\omega})$ triggers a switch in these state dynamics. IPA specifies how changes in $\boldsymbol{\theta}$ and $\boldsymbol{\omega}$ influence the state $\boldsymbol{s}(\boldsymbol{\theta}, \boldsymbol{\omega}, t, \boldsymbol{s}_0)$, as well

as event times $\tau_k(\boldsymbol{\theta}, \boldsymbol{\omega})$, $k = 1, 2, \ldots$, and, ultimately the cost function (22). We briefly review next the IPA framework for general stochastic hybrid systems as presented in [12].

Let $\{\tau_k(\theta)\}$, $k = 1, \ldots, K$, denote the occurrence times of all events in the state trajectory of a hybrid system with dynamics $\dot{x} = f_k(x, \theta, t)$ over an interval $[\tau_k(\theta), \tau_{k+1}(\theta))$, where $\theta \in \Theta$ is some parameter vector and Θ is a given compact, convex set. For convenience, we set $\tau_0 = 0$ and $\tau_{K+1} = T$. We use the Jacobian matrix notation: $x'(t) \equiv \frac{\partial x(\theta,t)}{\partial \theta}$ and $\tau_k' \equiv \frac{\partial \tau_k(\theta)}{\partial \theta}$, for all state and event time derivatives. It is shown in [12] that

$$\frac{d}{dt}x'(t) = \frac{\partial f_k(t)}{\partial x}x'(t) + \frac{\partial f_k(t)}{\partial \theta} \tag{23}$$

for $t \in [\tau_k, \tau_{k+1})$ with boundary condition:

$$x'(\tau_k^+) = x'(\tau_k^-) + [f_{k-1}(\tau_k^-) - f_k(\tau_k^+)]\tau_k' \tag{24}$$

for $k = 0, ...K$. In order to complete the evaluation of $x'(\tau_k^+)$ in (24), we need to determine τ_k'. If the event at τ_k is *exogenous*, $\tau_k' = 0$. However, if the event is *endogenous*, there exists a continuously differentiable function $g_k : \mathbb{R}^n \times \Theta \to \mathbb{R}$ such that $\tau_k = \min\{t > \tau_{k-1} \ : \ g_k(x(\theta, t), \theta) = 0\}$ and

$$\tau_k' = -[\frac{\partial g_k}{\partial x}f_k(\tau_k^-)]^{-1}(\frac{\partial g_k}{\partial \theta} + \frac{\partial g_k}{\partial x}x'(\tau_k^-)) \tag{25}$$

as long as $\frac{\partial g_k}{\partial x}f_k(\tau_k^-) \neq 0$. (details may be found in [12]).

Denote the time-varying cost along a given trajectory as $L(x, \theta, t)$, so the cost in the k-th interevent interval is $J_k(x, \theta) = \int_{\tau_k}^{\tau_{k+1}} L(x, \theta, t)dt$ and the total cost is $J(x, \theta) = \sum_{k=0}^{K} J_k(x, \theta)$. Differentiating and applying the Leibnitz rule with the observation that all terms of the form $L(x(\tau_k), \theta, \tau_k)\tau_k'$ are mutually canceled with $\tau_0 = 0, \tau_{K+1} = T$ fixed, we obtain

$$\begin{aligned}\frac{\partial J(x,\theta)}{\partial \theta} &= \sum_{k=0}^{K}\frac{\partial}{\partial \theta}\int_{\tau_k}^{\tau_{k+1}} L(x, \theta, t)dt \\ &= \sum_{k=0}^{K}\int_{\tau_k}^{\tau_{k+1}} \frac{\partial L(x,\theta,t)}{\partial x}x'(t) + \frac{\partial L(x,\theta,t)}{\partial \theta}dt\end{aligned} \tag{26}$$

In our setting, we have $L(x, \theta, t) = \sum_{i=1}^{M} R_i(t)$ from (22), which is not an explicit function of the state $\mathbf{x}(t) = [R_1(t), ...R_M(t), s_1(t)...s_N(t)]$. Thus, the gradient $\nabla J(\boldsymbol{\theta}, \boldsymbol{\omega}) = [\frac{\partial J(\boldsymbol{\theta},\boldsymbol{\omega})}{\partial \boldsymbol{\theta}}, \frac{\partial J(\boldsymbol{\theta},\boldsymbol{\omega})}{\partial \boldsymbol{\omega}}]^{\mathrm{T}}$ reduces to

$$\nabla J(\boldsymbol{\theta}, \boldsymbol{\omega}) = \frac{1}{T}\sum_{k=0}^{K}\sum_{i=1}^{M}\int_{\tau_k(\boldsymbol{\theta},\boldsymbol{\omega})}^{\tau_{k+1}(\boldsymbol{\theta},\boldsymbol{\omega})} \nabla R_i(t)dt \tag{27}$$

where $\nabla R_i(t) = [\frac{\partial R_i(t)}{\partial \boldsymbol{\theta}}, \frac{\partial R_i(t)}{\partial \boldsymbol{\omega}}]^{\mathrm{T}}$.

Applying (23)(24)(25), we can evaluate $\nabla R_i(t)$. In contrast to [3], in our problem agents are allowed to dwell on every target and IPA will optimize these dwelling times. Therefore, we need to consider all possible forms of control sequences: (i) $\pm 1 \to 0$, (ii) $0 \to \pm 1$, and (iii) $\pm 1 \to \mp 1$. We can then obtain from (5):

$$\frac{\partial R_i(t)}{\partial \theta_{j\xi}} = \frac{\partial R_i(\tau_k^+)}{\partial \theta_{j\xi}} - \begin{cases} 0 & \text{if } R_i(t) = 0, A_i < B_i P_i(\mathbf{s}(t)) \\ G\frac{\partial s_j(\tau_k^+)}{\partial \theta_{j\xi}}(t - \tau_k) & \text{otherwise} \end{cases} \tag{28}$$

$$\frac{\partial R_i(t)}{\partial \omega_{j\xi}} = \frac{\partial R_i(\tau_k^+)}{\partial \omega_{j\xi}} - \begin{cases} 0 & \text{if } R_i(t) = 0, A_i < B_i P_i(\mathbf{s}(t)) \\ G\frac{\partial s_j(\tau_k^+)}{\partial \omega_{j\xi}}(t - \tau_k) & \text{otherwise} \end{cases} \tag{29}$$

where $G = B_i \prod_{d \neq j}\left[1 - p_i(s_d(t))\right]\frac{\partial p_i(s_j)}{\partial s_j}$ and $\frac{\partial p_i(s_j)}{s_j} = \pm\frac{1}{r_j}$.

First, let us consider the events that cause switches in $\dot{R}_i(t)$ in (5) at time τ_k. For these events, the dynamics of $s_j(t)$ are continuous so that $\nabla s_j(\tau_k^-) = \nabla s_j(\tau_k^+)$. For target i,

$$\nabla R_i(\tau_k^+) = \begin{cases} \nabla R_i(\tau_k^-) & \textbf{if } \dot{R}_i(\tau_k^-) = 0, \\ & \dot{R}_i(\tau_k^+) = A_i - B_i P_i(\boldsymbol{s}(\tau_k^+)). \\ 0 & \textbf{if } \dot{R}_i(\tau_k^-) = A_i - B_i P_i(\boldsymbol{s}(\tau_k^-)), \\ & \dot{R}_i(\tau_k^+) = 0. \end{cases} \tag{30}$$

Second, let us consider events that cause switches in $\dot{s}_j(t) = u_j(t)$ at time τ_k. For these events, the dynamics of $R_i(t)$ are continuous so that $\nabla R_i(\tau_k^-) = \nabla R_i(\tau_k^+)$. In order to evaluate (28) and (29), we need $\frac{\partial s_j(\tau_k^+)}{\partial \theta_{j\xi}}$ and $\frac{\partial s_j(\tau_k^+)}{\partial \omega_{j\xi}}$. Clearly, these are not affected by future events and we only have to consider the current and prior control switches. Let $\theta_{j\xi}$ and $\omega_{j\xi}$ be the current switching point and dwelling time. Again, applying (23)(24)(25), we have

Case 1: $u_j(\tau_k^-) = \pm 1, u_j(\tau_k^+) = 0$

$$\frac{\partial s_j}{\partial \theta_{jl}}(\tau_k^+) = \begin{cases} 1 & \text{if } l = \xi \\ 0 & \text{if } l < \xi \end{cases} \tag{31}$$

$$\frac{\partial s_j}{\partial \omega_{jl}}(\tau_k^+) = 0 \quad \text{for all } l \leq \xi \tag{32}$$

Case 2: $u_j(\tau_k^-) = 0, u_j(\tau_k^+) = \pm 1$

$$\frac{\partial s_j}{\partial \theta_{jl}}(\tau_k^+) = \begin{cases} \frac{\partial s_j}{\partial \theta_{jl}}(\tau_k^-) - u_j(\tau_k^+)sgn\big(\theta_{j\xi} - \theta_{j(\xi-1)}\big) \text{ if } l = \xi \\ \frac{\partial s_j}{\partial \theta_{jl}}(\tau_k^-) - u_j(\tau_k^+)\Big[sgn(\theta_{jl} - \theta_{j(l-1)}) \\ \qquad\qquad -sgn(\theta_{j(l+1)} - \theta_{jl})\Big] \text{ if } l < \xi \end{cases} \tag{33}$$

$$\frac{\partial s_j}{\partial \omega_{jl}}(\tau_k^+) = -u_j(\tau_k^+) \quad \text{for all } l \leq \xi \tag{34}$$

Case 3: $u_j(\tau_k^-) = \pm 1, u_j(\tau_k^+) = \mp 1$

$$\frac{\partial s_j}{\partial \theta_{jl}}(\tau_k^+) = \begin{cases} 2 & \text{if } l = \xi \\ -\frac{\partial s_j}{\partial \theta_{jl}}(\tau_k^-) & \text{if } l < \xi \end{cases} \tag{35}$$

Details of these derivations can be found in [3]. An important difference arises in *Case 2* above, where $\tau_k = |\theta_{j1} - a| + \omega_{j1} + ... + |\theta_{j\xi} - \theta_{j(\xi-1)}| + \omega_{j\xi}$. We eliminate the constraints on the switching location that $\theta_{j\xi} \leq \theta_{j(\xi-1)}$ if ξ is even and $\theta_{j\xi} \geq \theta_{j(\xi-1)}$ if ξ is odd.

The event excitation problem. Note that all derivative updates above are limited to events occurring at times $\tau_k(\boldsymbol{\theta}, \boldsymbol{\omega})$, $k = 1, 2, \ldots$. Thus, this approach is scalable in the number of events characterizing the hybrid system, not its state space. While this is a distinct advantage, it also involves a potential drawback. In particular, it assumes that the events involved in IPA updates are observable along a state trajectory. However, if the current trajectory never reaches the vicinity of any target so as to be able to sense it and affect the overall uncertainty cost function, then any small perturbation to the trajectory will have no effect on the cost. As a result, IPA will fail as illustrated in Fig. 2: here, the single agent trajectory $s_1(\boldsymbol{\theta}, \boldsymbol{\omega}, t)$ is limited to include no event. Thus, if a gradient-based procedure is initialized with such $s_1(\boldsymbol{\theta}, \boldsymbol{\omega}, t)$, no event involved in the evaluation of $\nabla R_i(t)$ is "excited" and the cost gradient remains zero.

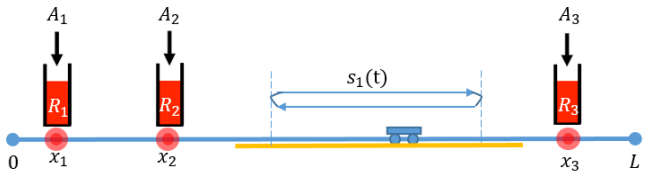

Fig. 2: An example of no event excitation leading to a failure of IPA finding an optimal agent trajectory. The yellow bar is the segment of the space covered by the agent.

In order to overcome this problem, we propose a modification of our cost metric by introducing a function $V(\cdot)$ with the property of "spreading" the value of some $R_i(t)$ over all points $w \in \Omega \equiv [0, L]$. Recalling Proposition 1, we limit ourselves to the subset $\mathcal{B} = [x_1, x_M] \subset \Omega$. Then, for all points $w \in \mathcal{B}$, we define $V(w, t)$ as a continuous density function which results in a total value equivalent to the weighted sum of the target values $\sum_{i=1}^{M} R_i(t)$. We impose the condition that $V(w, t)$ be monotonically decreasing in the Euclidean distance $\|w - x_i\|$. More precisely, we define $d_i^+(w) = \max\left(\|w - x_i\|, r\right)$ where $r = \min_{j=1,\ldots,N}\{r_j\}$ which ensures that $d_i^+(w) \geq r$. Thus, $d_i^+(w) = r > 0$ is fixed for all points within the target's vicinity, $w \in [x_i - r, x_i + r]$. We define

$$V(w, t) = \sum_{i=1}^{M} \frac{\alpha_i R_i(t)}{d_i^+(w)} \tag{36}$$

Note that $V(w, t)$ corresponds to the "total weighted reward density" at $w \in \mathcal{B}$. The weight α_i may be included to capture the relative importance of targets, but we shall henceforth set $\alpha_i = 1$ for all $i = 1, \ldots, M$ for simplicity. In order to differentiate points $w \in \mathcal{B}$ in terms of their location relative to the agents states $s_j(t)$, $j = 1, \ldots, N$, we also define the travel cost function

$$Q(w, \mathbf{s}(t)) = \sum_{j=1}^{N} \|s_j(t) - w\| \tag{37}$$

Using these definitions we introduce a new objective function component, which is added to the objective function in (6):

$$J_2(t) = \int_{\mathcal{B}} Q(w, \mathbf{s}(t)) V(w, t) dw \tag{38}$$

The significance of $J_2(t)$ is that it accounts for the movement of agents through $Q(w, \mathbf{s}(t))$ and captures the target state values through $V(w, t)$. Introducing this term in the objective function in the following creates a non-zero gradient even if the agent trajectories are not passing through any targets. We now define the metrics in (22) as $J_1(t)$ and incorporate the parametric $J_2(t)$ as an addition.

$$\min_{\boldsymbol{\theta} \in \Theta, \boldsymbol{\omega} \geq 0} J(\boldsymbol{\theta}, \boldsymbol{\omega}, T) = \frac{1}{T} \int_0^T \left[J_1(\boldsymbol{\theta}, \boldsymbol{\omega}, t) + e^{-\beta t} J_2(\boldsymbol{\theta}, \boldsymbol{\omega}, t)\right] dt \tag{39}$$

where $J_1(\boldsymbol{\theta}, \boldsymbol{\omega}, t) = \sum_{i=1}^{M} R_i(t)$ is the original uncertainty metric. This creates a continuous potential field for the agents which ensures a non-zero cost gradient even when the trajectories do not excite any events. This non-zero gradient will induce trajectory adjustments that naturally bring them toward ones with observable events. The inclusion of the factor $e^{-\beta t}$ with $\beta > 0$ is included so that as the number of IPA iterations increases, the effect of $J_2(\boldsymbol{\theta}, \boldsymbol{\omega}, t)$ is diminished and the original objective is ultimately recovered. The IPA derivative of $J_2(\boldsymbol{\theta}, \boldsymbol{\omega}, t)$ is

$$\frac{\partial}{\partial \boldsymbol{\theta}} \int_{\tau_k}^{\tau_{k+1}} \int_{\mathcal{B}} Q(w, \boldsymbol{\theta}, \boldsymbol{\omega}, \mathbf{s}(t), t) V(w, \boldsymbol{\theta}, \boldsymbol{\omega}, t) dw$$

$$= \int_{\tau_k}^{\tau_{k+1}} \int_{\mathcal{B}} \left[\frac{\partial Q(w, \boldsymbol{\theta}, \boldsymbol{\omega}, \mathbf{s}(t), t)}{\partial \boldsymbol{\theta}} V(w, \boldsymbol{\theta}, \boldsymbol{\omega}, t) \right. \tag{40}$$

$$\left. + Q(w, \boldsymbol{\theta}, \boldsymbol{\omega}, \mathbf{s}(t), t) \frac{\partial V(w, \boldsymbol{\theta}, \boldsymbol{\omega}, t)}{\partial \boldsymbol{\theta}}\right] dw \tag{41}$$

where the derivatives of $Q(w, \boldsymbol{\theta}, \boldsymbol{\omega}, \mathbf{s}(t), t)$ and $V(w, \boldsymbol{\theta}, \boldsymbol{\omega}, t)$ are obtained following the same procedure described previously. Before making this modification, the lack of event excitation in a state trajectory results in the total derivative (27) being zero. On the other hand, in (41) we observe that if no events occur, the second part in the integral, which involves $\frac{\partial V(\cdot)}{\partial \boldsymbol{\theta}}$ is zero, since $\sum_{i=1}^{M} \frac{\partial R_i(t)}{\partial \boldsymbol{\theta}} = 0$ at all t. However, the first part in the integral does not depend on events, but only the sensitivity of $Q(w, \boldsymbol{\theta}, \boldsymbol{\omega}, \mathbf{s}(t), t)$ in (37) with respect to the parameters $\boldsymbol{\theta}, \boldsymbol{\omega}$. As a result, agent trajectories are adjusted so as to eventually excite desired events and any gradient-based procedure we use in conjunction with IPA is no longer limited by the absence of event excitation.

IPA robustness to uncertainty modeling. Observe that the evaluation of $\nabla R_i(t)$, hence $\nabla J(\boldsymbol{\theta}, \boldsymbol{\omega})$, is *independent* of A_i, $i = 1, \ldots, M$, i.e., the parameters in our uncertainty model. In fact, the dependence of $\nabla R_i(t)$ on A_i, $i = 1, \ldots, M$, manifests itself through the event times τ_k, $k = 1, \ldots, K$, that do affect this evaluation, but they, unlike A_i which may be unknown, are directly observable during the gradient evaluation process. Thus, the IPA approach possesses an inherent *robustness* property: there is no need to explicitly model how uncertainty affects $R_i(t)$ in (5). Consequently, we may treat A_i as unknown without affecting

the solution approach (the values of $\nabla R_i(t)$ are obviously affected). We may also allow this uncertainty to be modeled through random processes $\{A_i(t)\}$, $i = 1, \ldots, M$. Under mild technical conditions on the statistical characteristics of $\{A_i(t)\}$ [12], the resulting $\nabla J(\boldsymbol{\theta}, \boldsymbol{\omega})$ is an unbiased estimate of a stochastic gradient.

V. Graph-based scheduling method

While the IPA-driven gradient-based approach described in Sec. IV offers several compelling advantages, it is not guaranteed to find a global optimum. In addition, it has been shown that in mission spaces of dimension greater than one, optimal trajectories cannot be described parametrically [4]. This motivates the use of an alternative approach where the targets are viewed as discrete tasks, leading naturally to a graph-based description of the problem [6], [19]–[22]. This higher level of abstraction allows one to guarantee an optimal solution, though at the cost of a significant increase in computational complexity. It is worth highlighting, however, that the complexity of such schemes is driven by the size of the graph and they are thus essentially invariant to the underlying dimensionality of the mission space.

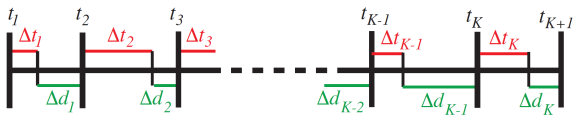

Fig. 3: Time sequence of a single agent on a given trajectory. The t_i are the time points where the agent begins to move to the next target in the sequence. Each move takes Δt_i units of time followed by a dwell period of Δd_i units of time during which information is collected from the target.

As illustrated in Fig. 3, our approach to the discrete setting is to divide the overall planning time horizon T for agent j into a sum of K_j consecutive time steps $\{t_j^1, t_j^2, ..., t_j^{K_j}\}$, $j = 1, \ldots, N$, with $t_j^1 = 0$. The dependence on j indicates that each agent may have a different discretization. We denote the end of the K-th step as $t_j^{K+1} = T$. Each step $k \in \{1, ..., K_j\}$ begins with a travel stage where the agent moves to a particular target i. Under the assumption that during the transition between targets each agent moves at its maximum speed of $|u_j| = 1$, the travel time is

$$\Delta t_j^k = |s_j^k(t_j^k) - x_i|. \tag{42}$$

Upon arriving at a target, the agent dwells for a time Δd_j^k. Note that due to the range-based nature of the sensing, the uncertainty actually begins to decrease before the arrival of the agent at the target and continues to decrease after the agent has departed until the target is out of the sensing range.

The problem of optimizing the u_j to minimize the average uncertainty over all the targets has been translated into a mixed integer programming (MIP) problem to select the sequence of targets and the dwell time at each target. Letting a_{ji}^k be a binary variable denoting whether agent j is assigned to target i at time step k, this MIP is

$$\min_{a_{ji}^k, \Delta d_j^k} J = \frac{1}{T} \sum_{i=1}^{M} \int_0^T R_i(t) dt \tag{43}$$

$$s.t. \quad a_{ji}^k \in \{1, 0\}, \quad \sum_{i=1}^{M} a_{ji}^k = 1, \quad \forall j, k \tag{44}$$

$$\sum_{k=1}^{K} \Delta t_j^k + \Delta d_j^k \leq T, \quad \forall j. \tag{45}$$

Note that we assume that each agent is assigned to a maximum of only one target at any one time. The IPA-driven approach has no such restriction. We break the solution of this problem into three parts: enumeration of all feasible trajectories, calculation of the cost of the feasible trajectories, and then selection of the optimal trajectory based on those costs. We focus on the case of a single agent for simplicity of description before generalizing to the multiple agent case.

The first part, namely determining feasible trajectories, is straightforward. Given the fixed time horizon T, the target locations, the locations of the agent at the start of the time horizon, and the maximum speed of the agent, a feasible trajectory is one where the sequence of targets can all be visited within the time horizon. Similarly, the third part simply involves comparing the trajectories and selecting the one with the minimal cost.

In the second part, the cost of each feasible trajectory must be determined. Suppose we have a given feasible trajectory with K targets in its sequence. Note that because a trajectory may include multiple visits to the same target, K may be larger than m (and may be much larger for large time horizons and small m). Let $\{i_1, i_2, \ldots, i_K\}$ denote the indices of the targets in the sequence. From (43), the cost of this trajectory is given by the optimization problem

$$\min_{\Delta d_j^k} J = \frac{1}{T} \sum_{i=1}^{M} \int_0^T R_i(t) dt$$

$$s.t. \quad \sum_{k=1}^{K} \Delta t^k + \Delta d^k \leq T.$$

Our approach to solving this optimization problem is to setup a recursive calculation. As illustrated in Fig. 3, since the travel times Δt_i are completely determined by the sequence alone, optimizing over the dwell times is equivalent to optimizing the switching times t_i. Assume for the moment that the switching times through t_{K-1} have been determined (and thus the first $K-2$ dwell times, $\Delta d^1, \ldots, \Delta d^{K-2}$ are known). The two final dwell times are completely determined by selecting the time t_K at which to switch the agent from target i_{K-1} to target i_K. This then gives us a simple single variable optimization problem

$$\min_{\Delta T_K} J = \frac{1}{\Delta T} \int_{t^{K-1}}^{T} \left(R_{i_{K-1}}(t) + R_{i_K}(t) \right) dt$$

where $\Delta T = T - t_{K-1}$. This allows the final switching time to be expressed as a function of the previous time

$t_K = t_K(t_{K-1})$. Repeating this leads to an expression of the optimal switching times as a nested sequence of optimization functions which can be solved numerically.

This same optimization procedure can be generalized to the case of multiple agents. The primary challenge is that the set of feasible trajectories, and the calculation of the cost of those trajectories, quickly becomes intractable since all possible combinations of assignments of multiple agents must be considered. The computational complexity can be mitigated somewhat by taking advantage of known properties of optimal solutions (as described in Sec. III).

Since the computationally complexity is exponential in the length of the time horizon, this approach is most feasible over short horizons. In prior work on linear systems, it was shown that an appropriately defined periodic schedule is sufficient to ensure the entire system remains controllable [23], [24]. In the current context, this translates to being able to keep the uncertainty of each of the targets arbitrarily close to zero. Motivated by this, we typically apply our discrete approach over a relatively short time horizon. If the resulting optimal trajectory is periodic, we extend it to longer horizons by simply repeating it.

VI. Simulation Examples

To demonstrate the performance of the gradient-based algorithm using the IPA scheme described in Sec. IV, we present two sets of numerical examples. The first set uses deterministic target locations and dynamics. The results are compared against the optimal found by the discrete scheduling algorithm of Sec. V. The second set demonstrates the robustness of the IPA scheme with respect to a stochastic uncertainty model.

The first simulation consists of a single agent performing a persistent monitoring task on three targets over a time horizon of 100 seconds. The targets are located at positions $x_1 = 5$, $x_2 = 10$, $x_3 = 15$ and their uncertainty dynamics in (5) are defined by the parameters $A_i = 1$, $B_i = 5$, and $R_i(0) = 1$ for all i. The agent has a sensing range of 2 and is initialized with $s(0) = 0$, $u(0) = 1$. The results from the IPA gradient descent approach are shown in Fig. 4. The top image shows the optimal trajectory of the agent determined after 1000 iterations of the IPA gradient descent while the bottom shows the evolution of the overall cost as a function of iteration number. The agent is moving through a periodic cycle of $x_1 \to x_2 \to x_3 \to x_2 \to x_1 \cdots$, dwelling for a short time at each target before moving to the next. Notice that the agent dwells for a shorter time at the center target since it visits that location twice per cycle. The second image in the figure shows that gradient descent converges within the first 100 iterations. This simulation aims to test the event driven IPA scheme with the discrete scheduling algorithm which yields optimal but suffers from computational intensity. Thus, we start with a short time horizon $T = 100$s. Event-driven IPA optimizes the trajectory fast but the convergence is somewhat unstable due to the lack of events within a short time horizon. The final cost is 26.11. The bottom images in Fig. 4 show the evolution of the target uncertainties.

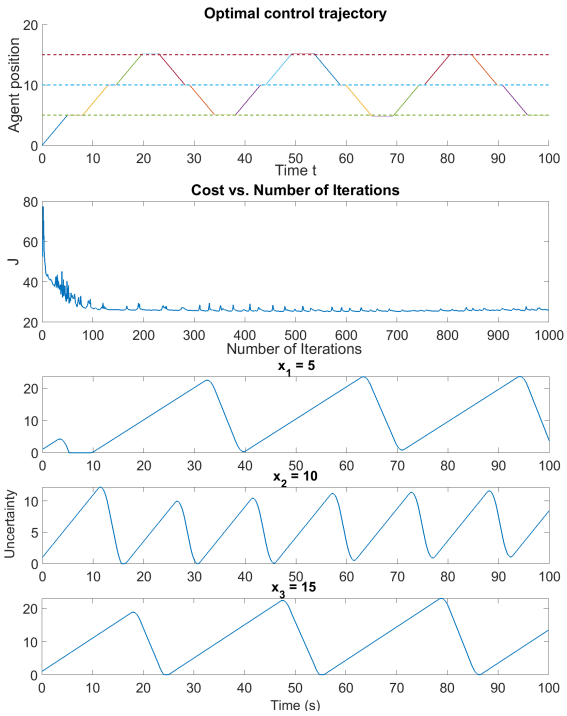

Fig. 4: A single agent monitoring three targets using the IPA-driven gradient descent algorithm. (top image) Agent trajectory. (second image) Calculated cost as a function of iteration in the gradient descent. The final cost is 26.11. (bottom images) Target uncertainties along the trajectory.

The corresponding result based on the discrete setting of Sec. V is essentially the same with the agent moving through the three targets in a periodic fashion as shown in Fig. 5. The only deviation from the IPA scheme occurs at the end of the horizon where the discrete approach returns to the center target. The final cost was 25.07, matching that of the IPA approach and thus verifying the approximate optimality of the solution found in Fig. 4.

The next simulation involves two agents and five targets over a time horizon of 500 seconds. The targets are located at $x_1 = 5$, $x_2 = 7$, $x_3 = 9$, $x_4 = 13$, $x_5 = 15$. The uncertain dynamics were the same as in the single agent, three target case. As before, the agents have a sensing range of 2 and are initialized at $s_1(0) = s_2(0) = 0$, with $u_1(0) = u_2(0) = 1$. The results from the event-driven IPA gradient descent approach are shown in Fig. 6. The solution is again periodic with the agents dividing the targets into two groups. Notice that the single agent on targets x_4 and x_5 is able to keep the uncertainties very close to zero since the targets are quite close relative to the sensing range of the agent. The other agent is able to hold its middle target (x_2) close to zero since it is visited more often. The uncertainties of targets x_1 and x_3 rise and decrease to zero constantly. The corresponding result based on the discrete setting is shown in Fig. 7. Rather than solve over the full horizon, the problem was solved over a 60 second horizon and then the periodic trajectory repeated to fill the 500 second horizon. The results are again very close to the event-driven IPA method.

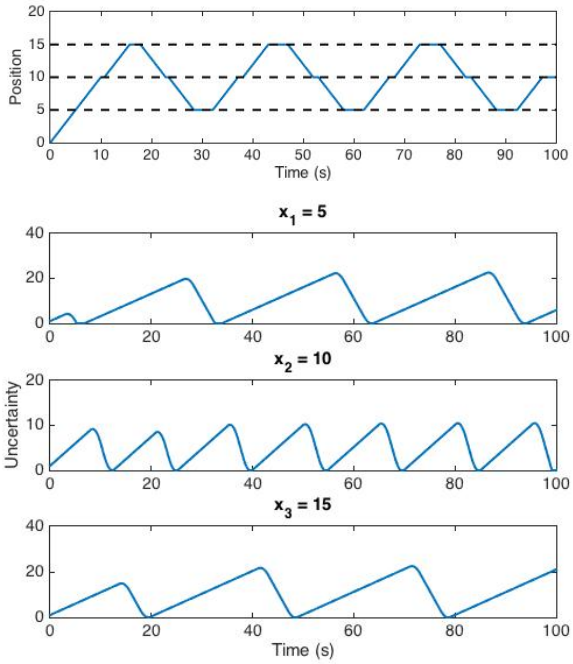

Fig. 5: A single agent monitoring three targets using the optimal discrete assignment and dwelling time. The final cost is 25.07. (top image) The agent trajectory is almost the same as in Fig. 4. (bottom images) Target uncertainties along the trajectory.

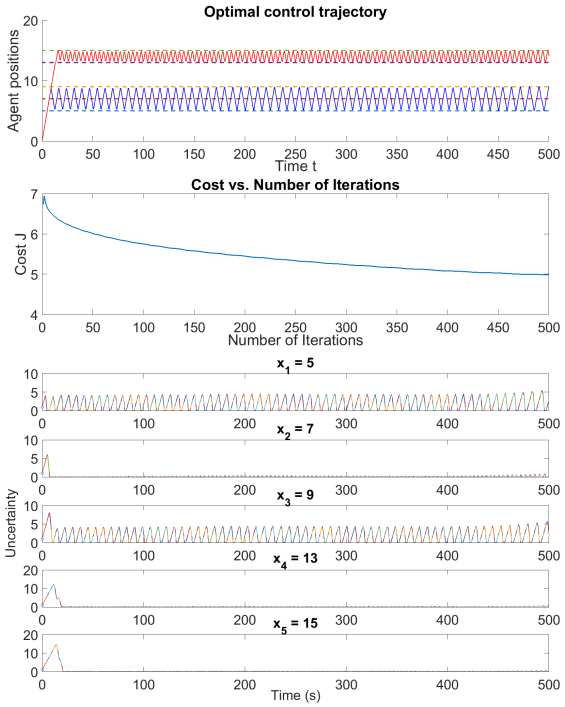

Fig. 6: Two agents monitoring five targets using the IPA gradient descent algorithm. (top image) Agent trajectories. (second image) Calculated cost as a function of iteration. The final cost is 4.99. (bottom images) Target uncertainty values along the above trajectories.

Note that the optimal trajectories in both one and two agent examples are bounded between $[5, 15]$ (positions of the first and last target), which is consistent with Proposition 1.

As mentioned earlier, the IPA robustness property allows us to handle stochastic uncertainty models at targets. We show a one-agent example in Fig. 8b where the uncertainty inflow rate $A_i(t)$ is uniformly distributed between $[0, 2]$ for all targets. In Fig. 8c, we introduce randomness by allowing target positions to vary uniformly over $[x_i - 0.25, x_i + 0.25]$. In both cases, the optimal cost in the stochastic models in Figs. 8b and 8c are close to the optimal cost of the deterministic case Fig. 8a where the parameter A_i and target positions x_i are the means of the associated random processes in the stochastic models. As expected, the convergence depends on the variance of these random processes.

The event excitation issue is addressed in Fig. 9a, where the agent trajectory is initialized so that it is not close to any of the targets. Using the original problem formulation (without the inclusion of $J_2(\boldsymbol{\theta}, \boldsymbol{\omega}, t)$ in (39), the initial trajectory and cost remain unchanged. After adding $J_2(\boldsymbol{\theta}, \boldsymbol{\omega}, t)$, the blue, green, and red curves in Fig. 9c show the trajectory adjustment after 5, 10, and 15 iterations respectively. After 100 iterations, the cost converges to 30.24 as shown in Fig. 9b which is close to the optimal cost in Fig. 8a where the target dynamics are the same.

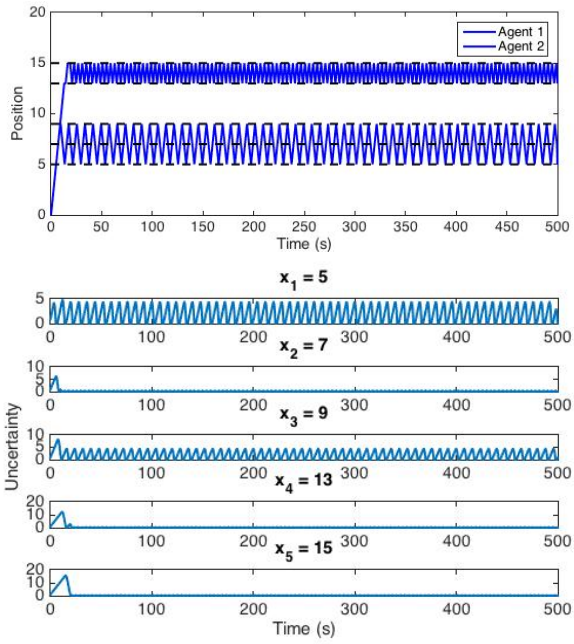

Fig. 7: Two agents monitoring five targets using the discrete assignment and dwelling time. The final cost was 4.92. (top image) Agent trajectories. (bottom images) Target uncertainty values along the above trajectories.

VII. Conclusion

We have formulated a persistent monitoring problem with the objective of controlling the movement of multiple cooperating agents so as to minimize an uncertainty metric asso-

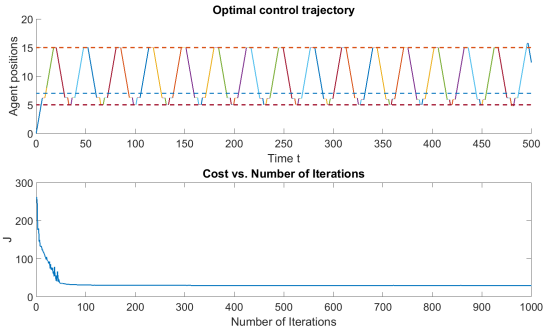

(a) Example of deterministic target model. Target positions $5, 7, 15$, dynamics parameter $A_i = 1, B = 5, r = 2$. $J^*(\boldsymbol{\theta}, \boldsymbol{\omega}) = 29.40$.

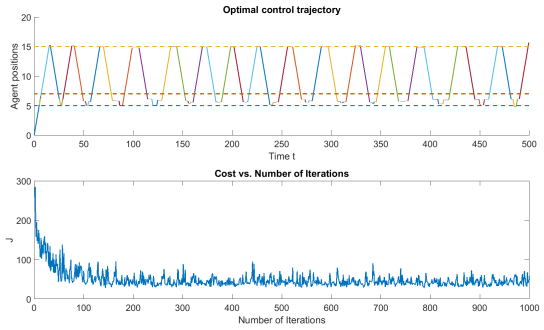

(b) Example with stochastic uncertainty processes. $A_i \sim U(0, 2)$. $J^*(\boldsymbol{\theta}, \boldsymbol{\omega}) = 42.46$.

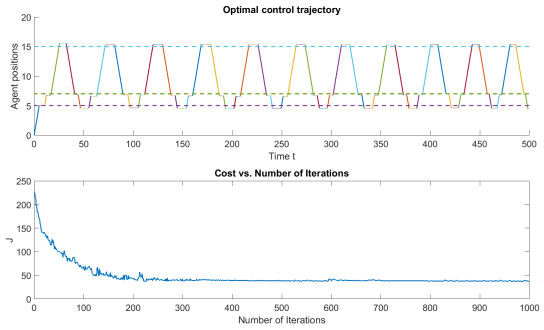

(c) Example with stochastic target locations $\sim U(x_i - 0.25, x_i + 0.25)$. $J^*(\boldsymbol{\theta}, \boldsymbol{\omega}) = 34.89$.

Fig. 8: Examples demonstrating IPA robustness with respect to stochastic uncertainty. (a)(b)(c) Top plot: optimal trajectory $s^*(t)$. Bottom plot: cost convergence.

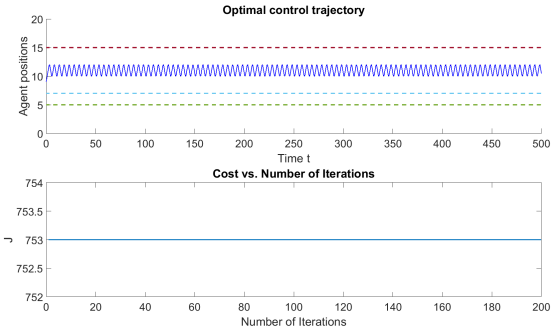

(a) A trajectory where IPA fails due to lack of event excitation. Top plot: agent trajectory. Bottom plot: cost convergence.

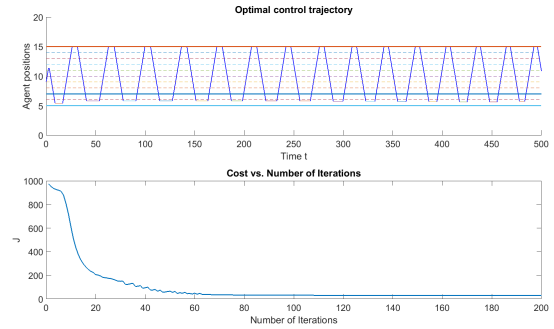

(b) IPA optimization after event excitation. Top plot: optimal agent trajectory. Bottom plot: cost convergence. $J^*(\boldsymbol{\theta}, \boldsymbol{\omega}) = 30.24$.

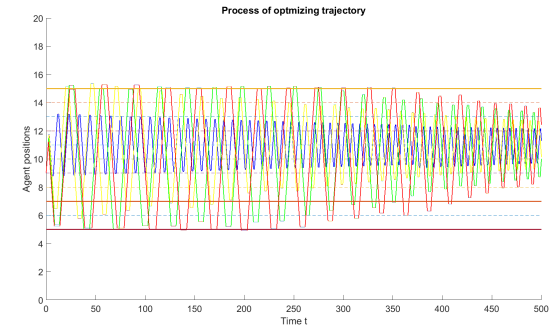

(c) Trajectory adjustments with event excitation after 5 (blue), 10 (green), and 15 (red) iterations.

Fig. 9: The event excitation issue. After adding $J_2(\boldsymbol{\theta}, \boldsymbol{\omega}, t)$, the trajectory adjusts to include targets, the cost converges to 30.24 which is close to the optimal cost in Fig.8a where the target dynamics are the same.

ciated with a finite number of targets. We have established properties of the optimal control solution which reduce the problem to a parametric optimization one. A complete on-line solution is given by Infinitesimal Perturbation Analysis (IPA) to evaluate the gradient of the objective function with respect to all parameters. We also address the case when IPA gradient estimation fails because of the lack of event excitation. We solve this problem by proposing a new metric for the objective function which creates a potential field guaranteeing that gradient values are non-zero. This approach is compared to an alternative graph-based task scheduling algorithm for determining an optimal sequence of target visits. Ongoing research is investigating how to extend these methodologies to higher dimensional mission spaces.

REFERENCES

[1] M. Zhong and C. G. Cassandras, "Distributed coverage control and data collection with mobile sensor networks," *Automatic Control, IEEE Transactions on*, vol. 56, no. 10, pp. 2445–2455, 2011.

[2] X. Sun and C. G. Cassandras, "Optimal dynamic formation control of multi-agent systems in environments with obstacles," *arXiv preprint arXiv:1508.04727*, 2015.

[3] C. Cassandras, X. Lin, and X. Ding, "An optimal control approach to the multi-agent persistent monitoring problem," *IEEE Transactions on Automatic Control*, vol. 58, no. 4, pp. 947–961, 2013.

[4] X. Lin and C. Cassandras, "An optimal control approach to the multi-agent persistent monitoring problem in two-dimensional spaces," in *Proc. of the IEEE Conference on Decision and Control*. IEEE, 2013, pp. 6886–6891.

[5] N. Michael, E. Stump, and K. Mohta, "Persistent surveillance with a team of mavs," in *2011 IEEE/RSJ International Conference on Intelligent Robots and Systems*, 2011.

[6] S. L. Smith, M. Schwager, and D. Rus, "Persistent monitoring of changing environments using a robot with limited range sensing," in *Proc. of the IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2011, pp. 5448–5455.

[7] Z. Shen and S. B. Andersson, "Tracking Nanometer-Scale Fluorescent Particles in Two Dimensions With a Confocal Microscope," *IEEE Transactions on Control Systems Technology*, vol. 19, no. 5, pp. 1269–1278, Sep. 2011.

[8] S. M. Cromer Berman, P. Walczak, and J. W. Bulte, "Tracking stem cells using magnetic nanoparticles," *Wiley Interdisciplinary Reviews: Nanomedicine and Nanobiotechnology*, vol. 3, no. 4, pp. 343–355, 2011.

[9] B. Horling and V. Lesser, "A survey of multi-agent organizational paradigms," *The Knowledge Engineering Review*, vol. 19, no. 04, pp. 281–316, 2004.

[10] J. Yu, S. Karaman, and D. Rus, "Persistent monitoring of events with stochastic arrivals at multiple stations," *Robotics, IEEE Transactions on*, vol. 31, no. 3, pp. 521–535, 2015.

[11] E. Stump and N. Michael, "Multi-robot persistent surveillance planning as a vehicle routing problem," in *Automation Science and Engineering (CASE), 2011 IEEE Conference on*. IEEE, 2011, pp. 569–575.

[12] C. G. Cassandras, Y. Wardi, C. G. Panayiotou, and C. Yao, "Perturbation analysis and optimization of stochastic hybrid systems," *European Journal of Control*, vol. 16, no. 6, pp. 642–661, 2010.

[13] Y. Wardi, R. Adams, and B. Melamed, "A unified approach to infinitesimal perturbation analysis in stochastic flow models: the single-stage case," *Automatic Control, IEEE Transactions on*, vol. 55, no. 1, pp. 89–103, 2010.

[14] M. Schwager, D. Rus, and J.-J. Slotine, "Decentralized, adaptive coverage control for networked robots," *The International Journal of Robotics Research*, vol. 28, no. 3, pp. 357–375, 2009.

[15] M. Cao, A. S. Morse, C. Yu, B. Anderson, S. Dasgupta *et al.*, "Maintaining a directed, triangular formation of mobile autonomous agents," *Communications in Information and Systems*, vol. 11, no. 1, p. 1, 2011.

[16] K.-K. Oh and H.-S. Ahn, "Formation control and network localization via orientation alignment," *Automatic Control, IEEE Transactions on*, vol. 59, no. 2, pp. 540–545, 2014.

[17] Y. Khazaeni and C. G. Cassandras, "Event excitation for event-driven control and optimization of multi-agent systems," in *IEEE International Workshop on Discrete Event Systems(WODES)*. IEEE, 2016.

[18] A. E. Bryson, *Applied optimal control: optimization, estimation and control*. CRC Press, 1975.

[19] X. Lan and M. Schwager, "Planning periodic persistent monitoring trajectories for sensing robots in gaussian random fields," in *IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2013, pp. 2415–2420.

[20] M. Lahijanian, J. Wasniewski, S. B. Andersson, and C. Belta, "Motion planning and control from temporal logic specifications with probabilistic satisfaction guarantees," in *IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2010, pp. 3227–3232.

[21] S. L. Smith, J. Tumova, C. Belta, and D. Rus, "Optimal path planning for surveillance with temporal logic constraints," *The International Journal of Robotics Research*, p. 0278364911417911, 2011.

[22] N. Mathew, S. L. Smith, and S. L. Waslander, "A graph-based approach to multi-robot rendezvous for recharging in persistent tasks," in *Robotics and Automation (ICRA), 2013 IEEE International Conference on*. IEEE, 2013, pp. 3497–3502.

[23] X. Yu and S. B. Andersson, "Effect of switching delay on a networked control system," in *Proc.of the IEEE Conference on Decision and Control (CDC)*. IEEE, 2013, pp. 5945–5950.

[24] X. Yu and S. B. Andersson, "Preservation of system properties for networked linear, time-invariant control systems in the presence of switching delays," in *Proc. of the IEEE Conference on Decision and Control*. IEEE, 2014, pp. 5260–5265.